arXiv:1402.6096v1 [cs.CG] 25 Feb 2014

Bounded-Angle Spanning Tree: Modeling Networks with Angular Constraints*

Rom Aschner Matthew J. Katz

Department of Computer Science, Ben-Gurion University, Israel
{romas,matya}@cs.bgu.ac.il

February 26, 2014

Abstract

We introduce a new structure for a set of points in the plane and an angle α, which is similar in flavor to a bounded-degree MST. We name this structure α-MST. Let P be a set of points in the plane and let $0 < \alpha \leq 2\pi$ be an angle. An α-ST of P is a spanning tree of the complete Euclidean graph induced by P, with the additional property that for each point $p \in P$, the smallest angle around p containing all the edges adjacent to p is at most α. An α-MST of P is then an α-ST of P of minimum weight. For $\alpha < \pi/3$, an α-ST does not always exist, and, for $\alpha \geq \pi/3$, it always exists [1, 2, 9]. In this paper, we study the problem of computing an α-MST for several common values of α.

Motivated by wireless networks, we formulate the problem in terms of directional antennas. With each point $p \in P$, we associate a wedge w_p of angle α and apex p. The goal is to assign an orientation and a radius r_p to each wedge w_p, such that the resulting graph is connected and its MST is an α-MST. (We draw an edge between p and q if $p \in \mathrm{w}_q$, $q \in \mathrm{w}_p$, and $|pq| \leq r_p, r_q$.) Unsurprisingly, the problem of computing an α-MST is NP-hard, at least for $\alpha = \pi$ and $\alpha = 2\pi/3$. We present constant-factor approximation algorithms for $\alpha = \pi/2, 2\pi/3, \pi$.

One of our major results is a surprising theorem for $\alpha = 2\pi/3$, which, besides being interesting from a geometric point of view, has important applications. For example, the theorem guarantees that given *any* set P of $3n$ points in the plane and *any* partitioning of the points into n triplets, one can orient the wedges of each triplet *independently*, such that the graph induced by P is connected. We apply the theorem to the *antenna conversion* problem.

1 Introduction

Let P be a set of points in the plane and let $0 < \alpha \leq 2\pi$ be an angle. An α-ST of P is a spanning tree of the complete Euclidean graph induced by P, with the additional property that for each point $p \in P$, the smallest angle around p containing all the edges adjacent to p is at most α. An α-MST of P is then an α-ST of P of minimum weight.

In this paper, we study the problem of computing an α-MST for several common values of α. For $\alpha < \pi/3$, an α-ST does not always exist (consider, e.g., an equilateral triangle). Moreover,

*Work by R. Aschner was partially supported by the Lynn and William Frankel Center for Computer Sciences. Work by M. Katz was partially supported by grant 1045/10 from the Israel Science Foundation. Work by M. Katz and R. Aschner was partially supported by grant 2010074 from the United States – Israel Binational Science Foundation.

it is well known that there always exists a Euclidean MST of degree at most 5. Therefore, it is interesting to focus on the range $\pi/3 \leq \alpha < 8\pi/5$.

Carmi et al. [9] showed that, for $\alpha = \pi/3$, an α-ST always exists. A somewhat simpler construction was subsequently proposed by Ackerman et al. [1]. Aichholzer et al. [2] have also obtained this result (together with additional related results), independently. However, in all these papers, the goal is to construct an α-ST (for $\alpha = \pi/3$) and not an α-MST.

The problem of computing an α-MST is similar in flavor to the problem of computing a Euclidean minimum weight degree-k spanning tree, which has been studied extensively (see, e.g., [4, 10,15,16,18]). A minimum weight degree-k spanning tree is a minimum weight spanning tree, such that the degree of each point is at most k, where the interesting values of k are 2,3, and 4. Notice that for $k = 2$ we get the Euclidean traveling salesman path problem.

The problem of computing an α-ST is closely related to problems in which one needs to compute a Hamiltonian path or cycle, with some restrictions on the angles. Fekete and Woeginger [13] showed that every set of points has a Hamiltonian *path*, such that all its angles are bounded by $\pi/2$. An alternative construction was given later in [9]. Fekete and Woeginger also conjectured that for every set of $2k \geq 8$ points there exists a Hamiltonian *cycle*, such that all its angles are bounded by $\pi/2$. Recently, Dumitrescu et al. [11] showed how to construct a Hamiltonian cycle whose angles are bounded by $2\pi/3$. As for lower bound, in [9] and, independently, in [11] it is shown that, for any $\varepsilon > 0$, there exists a set of points, for which any Hamiltonian path has an angle greater than $\pi/2 - \varepsilon$. The problem of finding Hamiltonian paths with large angles was also considered in [13], where it is conjectured that every point set admits a Hamiltonian path, whose angles are at least $\pi/6$; Bárány et al. [6] showed how to construct a path, whose angles are at least $\pi/9$.

Unsurprisingly, the problem of computing an α-MST is NP-hard, at least for $\alpha = \pi$ and $\alpha = 2\pi/3$. For $\alpha = \pi$, one can show this by a reduction from the problem of finding a Hamiltonian path in grid graphs of degree at most 3, which is known to be NP-hard [14]. The reduction is similar to the one described for the problem of computing a minimum weight degree-3 spanning tree [19], with a few simple adaptations. For $\alpha = 2\pi/3$, one can show this by a straight-forward reduction from Hamiltonian path in hexagonal grid graphs. Arkin et al. [3] showed that the problem of finding a Hamiltonian cycle in hexagonal grid graphs is NP-hard. However, with not too much effort, one can prove that finding a Hamiltonian path in hexagonal grid graphs is NP-hard as well.

Motivated by wireless networks, we formulate the problem of computing an α-MST in terms of directional antennas. In the last few years, directional antennas have received considerable attention (see, e.g., [7, 8, 17]), as they have some noticeable advantages over omni-directional antennas. In particular, they require less energy to reach a receiver at a given distance, and when broadcasting to this receiver the affected region is much smaller, reducing the probability of causing interference at friendly receivers or being subject to eves dropping by hostile receivers. With each point $p \in P$, we associate a wedge W_p of angle α and apex p. The goal now is to assign an orientation and a radius r_p to each wedge W_p, such that the resulting graph is connected and its MST is an α-MST. (We draw an edge between p and q if $p \in \mathrm{W}_q$, $q \in \mathrm{W}_p$, and $|pq| \leq r_p, r_q$.)

An interesting related problem is the *antenna conversion* problem. The *unit disk graph* of P, denoted $\mathrm{UDG}(P)$, is the graph in which there is an edge between p and q if $|pq| \leq 1$. This is the communication graph induced by P, where each point in P represents a transceiver equipped with an omni-directional antenna of radius 1. We assume that $\mathrm{UDG}(P)$ is connected. Suppose that one wishes to replace the omni-directional antennas with directional antennas of angle α. The goal now is to assign an orientation to each of the wedges W_p and to fix a common range $\delta = \delta(\alpha)$,

such that the resulting (symmetric) communication graph is a c-hop-spanner of $\text{UDG}(P)$, where $c = c(\alpha)$. Moreover, δ and c should be small constants. Aschner et al. [5] considered this problem for $\alpha = \pi/2$. Here we solve it for $\alpha = 2\pi/3$, using significantly smaller constants.

Our results. In Section 2 we focus on the case $\alpha = 2\pi/3$. We begin by describing a simple gadget: Given any set S of three points in the plane, we show how to orient the wedges associated with the points of S, such that G_S, the graph induced by S, is connected, and, moreover, the union of the wedges of S covers the plane. We then prove a surprising theorem, which, besides being interesting from a geometric point of view, has far-reaching applications, such as the one mentioned in the abstract. Informally, the theorem states that any two such gadgets are connected. That is, let S_1 and S_2 be two triplets of points in the plane, and assume that the wedges (associated with the points) of S_1 and, independently, of S_2 are oriented according to the gadget construction instructions, then the graph induced by $S_1 \cup S_2$ is connected. Proving this theorem turned out to be a very challenging task, due to the huge number of possible configurations that must be considered, and only after arriving at the current three-stage proof structure (see Section 2.2), were we able to complete the proof.

In Section 3, we present constant-factor approximation algorithms for computing an α-MST. In particular, we compute a 2-approximation for a π-MST, a 6-approximation for a $2\pi/3$-MST, and a 16-approximation for a $\pi/2$-MST. These approximations are actually with respect to a Euclidean MST, which is a lower bound for an α-MST, for any α. In Section 4, we present a solution to the antenna conversion problem for $\alpha = 2\pi/3$, based on the theorem above. Specifically, we construct, in $O(n \log n)$ time, a 6-hop-spanner of $\text{UDG}(P)$, in which each edge is of length at most 7. Finally, NP-hardness proofs for the problem of computing an α-MST, for $\alpha = \pi$ and $\alpha = 2\pi/3$, can be found in Section 5.

2 $\alpha = \frac{2\pi}{3}$

Notation. Let p be a point and let α be an angle. We denote the wedge of angle α and apex p by W_p. The left ray bounding W_p (when looking from p into W_p) is denoted by $\overleftarrow{\text{W}}_p$ and the right ray by $\overrightarrow{\text{W}}_p$. The bisector of W_p is denoted by $bis(\text{W}_p)$. The orientations of $\overleftarrow{\text{W}}_p$, $\overrightarrow{\text{W}}_p$, and $bis(\text{W}_p)$ are denoted by $\theta(\overleftarrow{\text{W}}_p)$, $\theta(\overrightarrow{\text{W}}_p)$, and $\theta(bis(\text{W}_p))$, respectively. The *orientation* of W_p is the orientation of its bisector and is denoted by $\theta(\text{W}_p)$. We denote the ray emanating from p of orientation $\theta(bis(\text{W}_p)) + 180$ by $\widetilde{\text{W}}_p$; its orientation is denoted by $\theta(\widetilde{\text{W}}_p)$.

Let S be a set of points, where each point $p \in S$ is associated with a wedge W_p of some orientation. The graph induced by S, denoted G_S, is the graph in which there is an edge between $p, q \in S$ if and only if $p \in \text{W}_q$ and $q \in \text{W}_p$. If there is an edge between p and q, we say that p and q are *connected* and denote this by $\{p\} \leftrightarrow \{q\}$. Similarly, if S_1 and S_2 are two such sets of points, and there exist a point p in S_1 and a point q in S_2 such that p and q are connected, then we say that S_1 and S_2 are *connected* and denote this by $\{S_1\} \leftrightarrow \{S_2\}$. The notation $\{p\} \not\leftrightarrow \{q\}$ means that p and q are not connected, and, similarly, $\{S_1\} \not\leftrightarrow \{S_2\}$ means that there does not exist a point in S_1 and a point in S_2 such that these points are connected.

2.1 The basic gadget

Claim 2.1. *Let $S = \{a, b, c\}$ be a set of three points in the plane, and set $\alpha = 2\pi/3$. Then, one can orient the wedges of S, such that G_S, the induced graph of S, contains a $2\pi/3$-ST of S, and the wedges of S cover the plane.*

Proof. Consider $\triangle abc$, and assume w.l.o.g. that $\angle b \leq \angle c \leq \angle a$. Then, $\angle b \leq 60$ and $\angle c < 90$. Draw $\triangle abc$, such that $\overline{bc}$ is horizontal (with b to the left of c) and a is not below the line containing $\overline{bc}$. Orient the wedges of S as follows (see Figure 1(a)): $\theta(\mathrm{w}_a) = 240$, $\theta(\mathrm{w}_b) = 0$, $\theta(\mathrm{w}_c) = 120$.

It is easy to see that the non-directed edges (a, b) and (b, c) are in the induced graph G_S. Thus, G_S contains a $2\pi/3$-ST. As for the second requirement, notice that w_a contains the wedge $\mathrm{w}_a{}'$ of orientation $\theta(\mathrm{w}_a)$ and apex b, and w_c contains the wedge $\mathrm{w}_c{}'$ of orientation $\theta(\mathrm{w}_c)$ and apex b. But, clearly, $\mathrm{w}_a{}' \cup \mathrm{w}_b \cup \mathrm{w}_c{}' = \mathbb{R}$. □

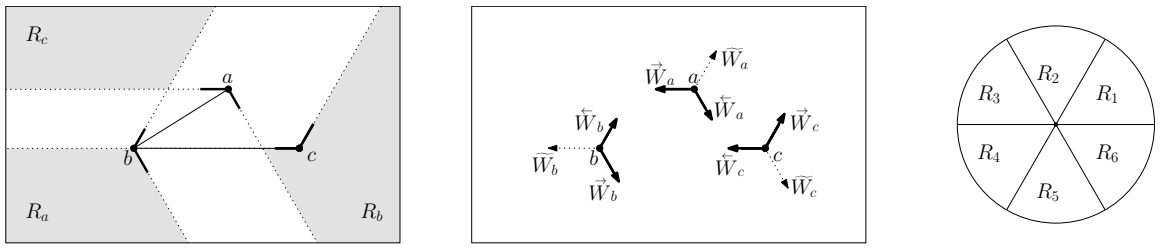

(a) (b) (c)

Figure 1: (a) The basic gadget of Claim 2.1. $\theta(\mathrm{w}_a) = 240$, $\theta(\mathrm{w}_b) = 0$, and $\theta(\mathrm{w}_c) = 120$. A point p is in region R_a if and only if $p \in \mathrm{w}_a$ and $p \notin \mathrm{w}_b, \mathrm{w}_c$, i.e., $R_a = \mathrm{w}_a \setminus (\mathrm{w}_b \cup \mathrm{w}_c)$. Regions R_b and R_c are defined analogously. (b) $\theta(\overleftarrow{\mathrm{w}}_a) = \theta(\overrightarrow{\mathrm{w}}_b) = \theta(\widetilde{\mathrm{w}}_c) = 300$, $\theta(\overrightarrow{\mathrm{w}}_a) = \theta(\widetilde{\mathrm{w}}_b) = \theta(\overleftarrow{\mathrm{w}}_c) = 180$, and $\theta(\widetilde{\mathrm{w}}_a) = \theta(\overleftarrow{\mathrm{w}}_b) = \theta(\overrightarrow{\mathrm{w}}_c) = 60$. (c) The six ranges $R_1, \ldots, R_6$.

The gadget of Claim 2.1 has some noticeable properties:

Property 1. For any $x \in S$, the orientations of the wedges of S are $\theta(\mathrm{w}_x)$ and $\theta(\mathrm{w}_x) \pm 120$.

Property 2. For any $x \in S$, the orientations of the rays bounding the wedges of S are $\theta(\mathrm{w}_x) \pm 60$ and $\theta(\mathrm{w}_x) + 180$. Moreover, each of these three orientations appears exactly twice, once as the orientation of a left ray bounding some wedge and once as the orientation of a right ray bounding some other wedge (see Figure 1(b)).

Property 3. Consider any two wedges w_x and w_y and the four rays defining them. Then, by Property 2, exactly two of these rays, ρ_1 from w_x and ρ_2 from w_y, have the same orientation. Let l be a line intersecting both ρ_1 and ρ_2 and perpendicular to ρ_1 (and to ρ_2). Then, $\mathrm{w}_x \cup \mathrm{w}_y$ covers the halfplane defined by l that does not include the points x and y.

Finally, let R_i denote the range $((i-1)60, i60)$, for $i = 1, \ldots, 6$ (see Figure 1(c)).

2.2 The induced graph of $S_1 \cup S_2$ is connected

In this section, we prove the following surprising theorem (Theorem 2.4), which, as mentioned, has far-reaching applications. Let $S_1 = \{a, b, c\}$ and S_2 be two triplets of points in the plane, and assume that the wedges (associated with the points) of S_1 and, independently, of S_2 are oriented according to the proof of Claim 2.1. Then, the induced graph of $S_1 \cup S_2$ is connected.

In order to cope with the huge number of cases, we prove Theorem 2.4 in three stages. In the first stage (Lemma 2.2), we prove the statement assuming that both induced graphs of S_1 and of S_2 are cliques. In the second stage (Lemma 2.3), we prove the statement assuming only one of the induced graphs is a clique, using, of course, Lemma 2.2. Finally, in the third stage (Theorem 2.4), we prove the statement without any additional assumptions, using Lemma 2.3.

Throughout this section, we assume (as in the proof of Claim 2.1) that, in Δabc, $\angle b \leq \angle c \leq \angle a$, $\overline{bc}$ is horizontal, with b to the left of c, and a is not below the line l containing $\overline{bc}$ (see Figure 1(a)).

Lemma 2.2 (Two cliques)**.** *Let $S_1 = \{a, b, c\}$ and S_2 be two triplets of points in the plane and let $\alpha = 2\pi/3$. Assume that the wedges (associated with the points) of S_1 and, independently, of S_2 are oriented according to the proof of Claim 2.1, and that both induced graphs, G_{S_1} and G_{S_2}, are cliques. Then, the induced graph $G_{S_1 \cup S_2}$ is connected.*

Proof. The wedges of S_2 cover the plane, in particular they cover all points of S_1. Therefore, we distinguish between three (not necessarily disjoint) cases: (i) there exists a point $x \in S_2$ such that w_x covers all points of S_1, (ii) there exists a point $x \in S_2$ such that w_x covers exactly two points of S_1, and (iii) the wedge of each point in S_2 covers exactly one point of S_1.

Case (i): There exists a point $x \in S_2$ such that w_x covers all points of S_1. Since the wedges of S_1 cover the plane, at least one of them must cover x, and therefore $\{x\} \leftrightarrow \{S_1\}$.

Case (ii): There exists a point $x \in S_2$ such that w_x covers exactly two points of S_1. We divide this case into three sub-cases, according to which two points of S_1 are covered by w_x.

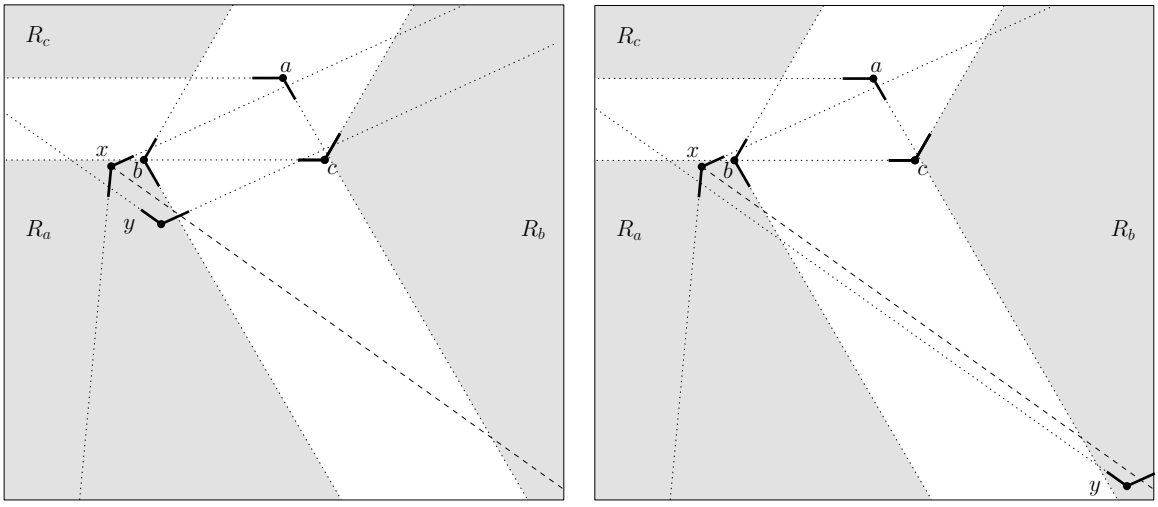

(a) $y \in R_a$ (b) $y \in \mathrm{w}_b$

Figure 2: Proof of Lemma 2.2, Case (ii)(1).

(1) w_x covers b and c and does not cover a. Assume $\{x\} \not\leftrightarrow \{b, c\}$ (since otherwise we are done), then $x \in R_a$ and one of the rays of w_x intersects $\overline{ab}$ and $\overline{ac}$. Notice that this ray must be $\overleftarrow{\mathrm{w}}_x$ and that $\overleftarrow{\mathrm{w}}_x$ also intersects $\overleftarrow{\mathrm{w}}_b$ (see Figure 2). Since x lies below l, $\overleftarrow{\mathrm{w}}_x$ intersects $\overleftarrow{\mathrm{w}}_b$, and $\theta(\overleftarrow{\mathrm{w}}_b) = 60$, we have that $\theta(\overleftarrow{\mathrm{w}}_x) \in R_1$. It follows that $\theta(\overrightarrow{\mathrm{w}}_x) \in R_5$, $\theta(\mathrm{w}_x) \in R_6$, and $\theta(\widetilde{\mathrm{w}}_x) \in R_3$. Therefore, $bis(\mathrm{w}_x)$ (whose orientation is $\theta(\mathrm{w}_x)$) does not intersect l. Let y be the point of S_2 such that $\theta(\overleftarrow{\mathrm{w}}_y) = \theta(\widetilde{\mathrm{w}}_x) \in R_3$ and $\theta(\overrightarrow{\mathrm{w}}_y) = \theta(\overleftarrow{\mathrm{w}}_x) \in R_1$. Since $\{x\} \leftrightarrow \{y\}$, we have that $y \in \mathrm{w}_x$ and

y lies to the right of $bis(\mathrm{w}_x)$. Notice that w_y contains the (imaginary) wedge of orientation $\theta(\mathrm{w}_y)$ and apex x. If $y \in R_a$ (see Figure 2(a)), then $\{y\} \leftrightarrow \{a\}$, since w_y covers a. Otherwise, $y \in \mathrm{w}_b$ and in particular y lies to the right of b (see Figure 2(b)). In this case we show that $\{y\} \leftrightarrow \{b\}$. Indeed, $\overrightarrow{\mathrm{w}}_y$ intersects l to the right of b, since $\theta(\overrightarrow{\mathrm{w}}_y) \in R_1$, and, since $l(\overleftarrow{\mathrm{w}}_y)$ is parallel to $l(bis(\mathrm{w}_x))$ and below it, we have that $\overleftarrow{\mathrm{w}}_y$ intersects l to the left of b. We conclude that $b \in \mathrm{w}_y$ and $\{y\} \leftrightarrow \{b\}$.

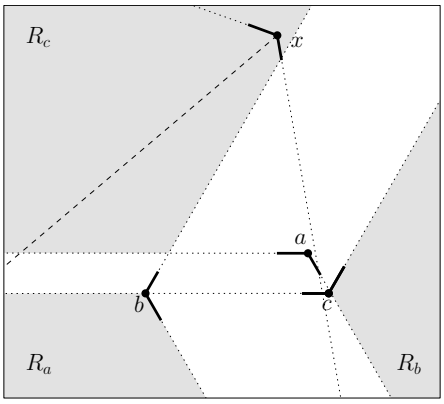

Figure 3: Proof of Lemma 2.2, Case (ii)(2).

(2) w_x covers a and b and does not cover c. Assume $\{x\} \not\leftrightarrow \{a, b\}$ (since otherwise we are done), then $x \in R_c$ and one of the rays of w_x intersects $\overline{ac}$ and $\overline{bc}$. Notice that this ray must be $\overleftarrow{\mathrm{w}}_x$ and that $\overleftarrow{\mathrm{w}}_x$ also intersects $\overleftarrow{\mathrm{w}}_a$ (see Figure 3). Since x lies above l, $\overleftarrow{\mathrm{w}}_x$ intersects $\overleftarrow{\mathrm{w}}_a$, and $\theta(\overleftarrow{\mathrm{w}}_a) = 300$, we have that $\theta(\overleftarrow{\mathrm{w}}_x) \in R_5$. It follows that $\theta(\overrightarrow{\mathrm{w}}_x) \in R_3$, $\theta(\mathrm{w}_x) \in R_4$, and $\theta(\widetilde{\mathrm{w}}_x) \in R_1$. The rest of the proof for this case is very similar to the proof of Case (ii)(1), thus we omit further details.

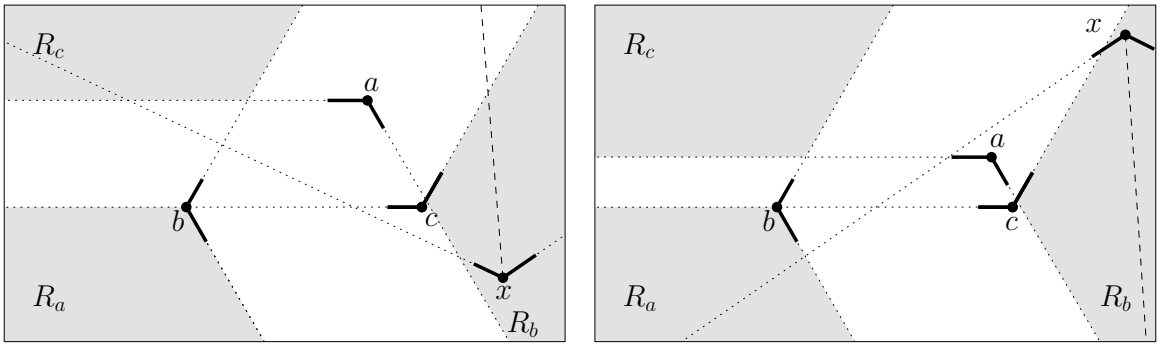

(a) $\overleftarrow{\mathrm{w}}_x$ intersects $\overline{ab}$ and $\overline{bc}$

(b) $\overrightarrow{\mathrm{w}}_x$ intersects $\overline{ab}$ and $\overline{bc}$

Figure 4: Proof of Lemma 2.2, Case (ii)(3).

(3) w_x covers a and c and does not cover b. Assume $\{x\} \not\leftrightarrow \{a, c\}$ (since otherwise we are done), then $x \in R_b$, and one of the rays of w_x intersects $\overline{ab}$ and $\overline{bc}$. Notice that this ray can be either $\overleftarrow{\mathrm{w}}_x$ or $\overrightarrow{\mathrm{w}}_x$.

If it is $\overleftarrow{\mathrm{w}}_x$ (see Figure 4(a)), then the orientations associated with w_x are: $\theta(\overleftarrow{\mathrm{w}}_x) \in R_3$, $\theta(\overrightarrow{\mathrm{w}}_x) \in R_1$, and $\theta(\widetilde{\mathrm{w}}_x) \in R_5$. The rest of the proof for this branch is very similar to the proof of Case (ii)(1), thus we omit further details.

If the ray intersecting $\overline{ab}$ and $\overline{bc}$ is $\overrightarrow{\mathrm{w}}_x$ (see Figure 4(b)), then the orientations associated with w_x are: $\theta(\overleftarrow{\mathrm{w}}_x) \in R_6$, $\theta(\overrightarrow{\mathrm{w}}_x) \in R_4$, and $\theta(\widetilde{\mathrm{w}}_x) \in R_2$. Again, the rest of the proof for this branch is very similar to the proof of Case (ii)(1), thus we omit further details.

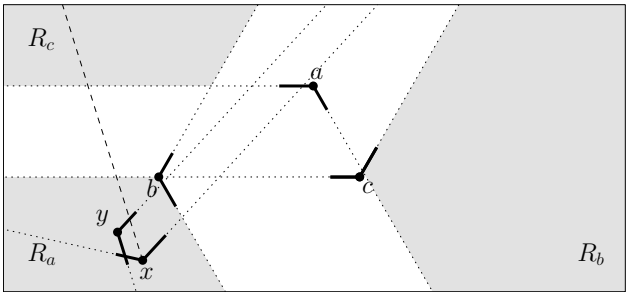

Figure 5: Proof of Lemma 2.2, Case (iii).

Case (iii): The wedge of each point in S_2 covers exactly one point of S_1. We may assume that this condition also holds for the wedges of S_1; that is, the wedge of each point in S_1 covers exactly one point of S_2. Since, otherwise, we can simply interchange the set names. It follows that each point of S_2 lies in its own private region among the regions R_a, R_b, and R_c.

Let x be the point that lies in R_a. We claim that $\{x\} \leftrightarrow \{a\}$. Assume that $\{x\} \not\leftrightarrow \{a\}$. We show that there exists a point $y \in S_2$ that covers two points of S_1. If w_x covers b (see Figure 5), then $\theta(\overrightarrow{\mathrm{w}}_x) \in (0, 120)$, which implies that $\theta(\widetilde{\mathrm{w}}_x) \in (240, 360)$. Let y be the point of S_2 such that $\theta(\overrightarrow{\mathrm{w}}_y) = \theta(\widetilde{\mathrm{w}}_x)$ and $\theta(\overleftarrow{\mathrm{w}}_y) = \theta(\overrightarrow{\mathrm{w}}_x)$. Since $\{y\} \leftrightarrow \{x\}$, we have that $y \in \mathrm{w}_x$ and y lies to the left of $bis(\mathrm{w}_x)$, but then w_y must cover a and c – contradiction. If w_x covers c, then $\theta(\overleftarrow{\mathrm{w}}_x) \in (0, 120)$, which implies that $\theta(\widetilde{\mathrm{w}}_x) \in (120, 240)$. Let y be the point of S_2 such that $\theta(\overleftarrow{\mathrm{w}}_y) = \theta(\widetilde{\mathrm{w}}_x)$ and $\theta(\overrightarrow{\mathrm{w}}_y) = \theta(\overleftarrow{\mathrm{w}}_x)$. Since $\{y\} \leftrightarrow \{x\}$, we have that $y \in \mathrm{w}_x$ and y lies to the right of $bis(\mathrm{w}_x)$, but then w_y must cover a and b – contradiction.

□

Lemma 2.3 (One clique)**.** *Let $S_1 = \{a, b, c\}$ and S_2 be two triplets of points in the plane and let $\alpha = 2\pi/3$. Assume that the wedges of S_1 and, independently, of S_2 are oriented according to the proof of Claim 2.1, and that the induced graph G_{S_2} is a clique. Then, the induced graph $G_{S_1 \cup S_2}$ is connected.*

Proof. If the induced graph G_{S_1} is also a clique, then, by Lemma 2.2, we are done. Assume therefore that G_{S_1} is not a clique. Let c' be the intersection point of $\overleftarrow{\mathrm{w}}_a$ and $\overleftarrow{\mathrm{w}}_c$ (see Figure 6), and consider the wedge $\mathrm{w}_{c'}$ of orientation $\theta(\mathrm{w}_{c'}) = \theta(\mathrm{w}_c)$ and apex c'. The graph induced by $\{a, b, c'\}$ is a clique, and therefore, by Lemma 2.2, $\{a, b, c'\} \leftrightarrow \{S_2\}$. If $\{a, b\} \leftrightarrow \{S_2\}$, then we are done, so assume that $\{c'\} \leftrightarrow \{S_2\}$. Let x be a point of S_2 such that $\{x\} \leftrightarrow \{c'\}$, and assume that w_x does not cover c (if it does, then $\{x\} \leftrightarrow \{c\}$, since $\mathrm{w}_{c'} \subseteq \mathrm{w}_c$). Then, x lies above l and $\overleftarrow{\mathrm{w}}_x$ intersects $\overline{cc'}$. Below we consider the three cases: (i) $\overrightarrow{\mathrm{w}}_x$ intersects $\overline{bc'}$, (ii) $\overrightarrow{\mathrm{w}}_x$ intersects l to the left of b, and (iii) $\overrightarrow{\mathrm{w}}_x$ does not intersect l. However, in the first case (i.e., Case (i)) and in sub-cases (1) and (2) of the second case (i.e., Case (ii)(1) and Case (ii)(2)) we refrain from using the assumption that G_{S_2} is a clique. This is because these cases appear again later in the proof of Theorem 2.4, where we may not assume that G_{S_2} is a clique.

Case (i): $\overrightarrow{\mathrm{w}}_x$ intersects $\overline{bc'}$ (see Figure 6(a)). Notice that in this case w_x does not cover points b and c. Since $\theta(\overleftarrow{\mathrm{w}}_x) < 360$ and $\theta(\overrightarrow{\mathrm{w}}_x) > 180$, we get that $\theta(\overleftarrow{\mathrm{w}}_x) \in R_6$, $\theta(\overrightarrow{\mathrm{w}}_x) \in R_4$, and $\theta(\widetilde{\mathrm{w}}_x) \in R_2$.

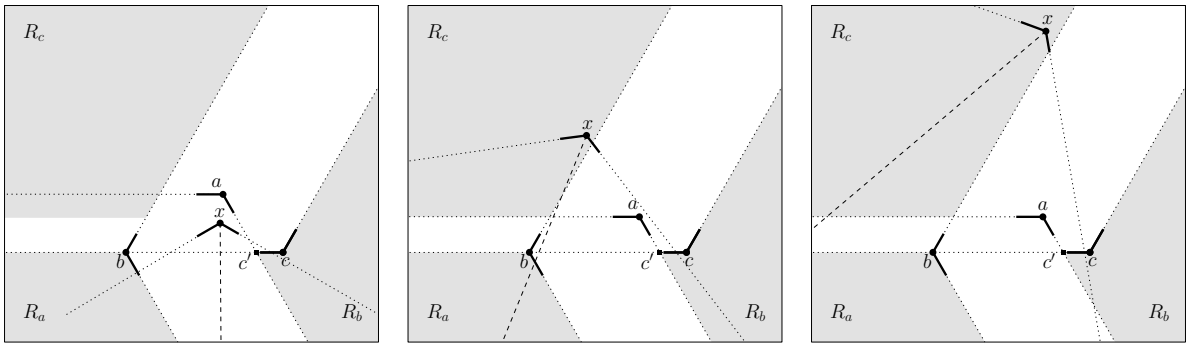

(a) Case (i) (b) Case (ii) (c) Case (iii)

Figure 6: Proof of Lemma 2.3.

Between the two points in $S_2 \setminus \{x\}$, let y be the one whose wedge covers more points of S_1; in case of tie, let y be any one of them. We know that one of w_y's rays has orientation in R_2. There are five sub-cases:

(1) w_y covers all points of S_1. There must exist a point in S_1 that covers y, so we are done.

(2) w_y covers b and c and does not cover a. If $\{y\} \leftrightarrow \{b, c\}$, then we are done. Otherwise, $y \in R_a$. Now, since $\theta(\overleftarrow{\mathrm{w}}_b) = 60$ and w_y must cover b and c and avoid a, we get that $\theta(\overleftarrow{\mathrm{w}}_y) \in R_1$. But, this is impossible, since R_1 is not among the three relevant ranges mentioned above.

(3) w_y covers a and b and does not cover c. If $\{y\} \leftrightarrow \{a, b\}$, then we are done. Otherwise, $y \in R_c$. We show that this is impossible. If $\theta(\overrightarrow{\mathrm{w}}_y) \in R_2$, then $\theta(\overleftarrow{\mathrm{w}}_y) \in R_4$, and $a, b \notin \mathrm{w}_y$. And, if $\theta(\overleftarrow{\mathrm{w}}_y) \in R_2$, then $\theta(\overrightarrow{\mathrm{w}}_y) \in R_6$, and w_y must also cover c.

(4) w_y covers a and c and does not cover b. This case is analogous to the previous one.

(5) w_y covers exactly one point of S_1. Therefore, the wedge of each point in S_2 covers exactly one point of S_1. Since w_x does not cover points b and c, it must cover a. Assume, w.l.o.g., that w_y covers c and w_z, the wedge of the remaining point, covers b. Next, we show that this is impossible. Indeed, if $\theta(\overrightarrow{\mathrm{w}}_y) \in R_2$ and $\theta(\overleftarrow{\mathrm{w}}_y) \in R_4$, then w_y must also cover a and b. And, if $\theta(\overrightarrow{\mathrm{w}}_z) \in R_2$ and $\theta(\overleftarrow{\mathrm{w}}_z) \in R_4$, then both y and z must lie below l. (Since, if y is above l, then $\{y\} \leftrightarrow \{c\}$, and, if z is above l, then $\{z\} \leftrightarrow \{b\}$). Therefore, $\mathrm{w}_y \cup \mathrm{w}_z$ covers the halfplane above l (see Property 3), and, in particular, at least one of the two wedges covers a.

Case (ii): $\overrightarrow{\mathrm{w}}_x$ intersects l to the left of b (see Figure 6(b)). In this case, as in Case (i), $\theta(\overleftarrow{\mathrm{w}}_x) \in R_6$, $\theta(\overrightarrow{\mathrm{w}}_x) \in R_4$, and $\theta(\widetilde{\mathrm{w}}_x) \in R_2$. Notice that in this case $b \in \mathrm{w}_x$, so we assume that $x \notin \mathrm{w}_b$, since otherwise $\{x\} \leftrightarrow \{b\}$. Let y be a point of S_2 whose wedge covers c. We distinguish between three sub-cases:

(1) w_y covers all points of S_1. There must exist a point in S_1 that covers y, so we are done.

(2) w_y covers exactly two points of S_1. If w_y covers b and c and $\{y\} \not\leftrightarrow \{b, c\}$, then $y \in R_a$ and either $\theta(\mathrm{w}_y) \in R_1$ or $\theta(\mathrm{w}_y) \in R_3$. However, in both cases, w_y must also cover a – contradiction. (Since, in the former case, $\overleftarrow{\mathrm{w}}_y$ does not intersect $\overleftarrow{\mathrm{w}}_b$, and in the latter case, $\overrightarrow{\mathrm{w}}_y$ does not intersect $\overleftarrow{\mathrm{w}}_a$.) If w_y covers a and c and $\{y\} \not\leftrightarrow \{a, c\}$, then $y \in R_b$ and $\theta(\mathrm{w}_y) \in R_3$. However, in the case, w_y must also cover b – contradiction. (Since $\overrightarrow{\mathrm{w}}_y$ passes above a and is directed upwards, and $\overleftarrow{\mathrm{w}}_y$ passes below c and is directed downward.)

(3) w_y covers exactly one point of S_2, namely, c. We know that either $\theta(\mathrm{w}_y) \in R_1$ or $\theta(\mathrm{w}_y) \in R_3$. In the latter case, w_y must also cover b, which is impossible. In the former case, if y is above l,

then $\{y\} \leftrightarrow \{c\}$, so y is necessarily below l. Let z be the remaining point. Then, $\theta(\mathrm{w}_z) \in R_3$. We show below that $\{z\} \leftrightarrow \{a, b\}$. Notice first that $\overleftarrow{\mathrm{w}}_y$ separates between a and c and between b and c, since $\theta(\mathrm{w}_y) \in R_1$ and w_y covers only c. Since G_{S_2} is a clique, we know that $\{y\} \leftrightarrow \{z\}$, and therefore z lies to the right of $bis(\mathrm{w}_y)$. Clearly, a and b lie to the left of $\overrightarrow{\mathrm{w}}_z$ (whose orientation is in R_2), and to the right of $\overleftarrow{\mathrm{w}}_z$ (whose orientation is in R_4). In other words, w_z covers both a and b. Notice also that $z \notin R_c$, since $bis(\mathrm{w}_y)$ (whose orientation is in R_1) intersects l to the right of b, and z lies to the right of $bis(\mathrm{w}_y)$. Therefore, either w_a or w_b (or both) covers z. We conclude that $\{z\} \leftrightarrow \{a, b\}$.

Case (iii): $\overrightarrow{\mathrm{w}}_x$ does not intersect l, i.e., $\theta(\overrightarrow{\mathrm{w}}_x) < 180$ (see Figure 6(c)). Since w_x covers b, we may assume that $x \notin \mathrm{w}_b$. Therefore, $\theta(\overleftarrow{\mathrm{w}}_x) > 240$. We thus have that $\theta(\overrightarrow{\mathrm{w}}_x) \in R_3$ and $\theta(\overleftarrow{\mathrm{w}}_x) \in R_5$. Notice that $bis(\mathrm{w}_x)$ (whose orientation is in R_4) intersects l to the right of b. Moreover, w_x necessarily covers a, since $\theta(\overleftarrow{\mathrm{w}}_x) \in R_5$ and $\overleftarrow{\mathrm{w}}_x$ intersects l between c' and c. Let y be the point of S_2 such that $\theta(\mathrm{w}_y) \in R_6$. Since G_{S_2} is a clique, we know that $\{x\} \leftrightarrow \{y\}$, and therefore y lies to the right of $bis(\mathrm{w}_x)$. If y is above l, then $\{y\} \leftrightarrow \{c\}$. Otherwise, y is below l and in w_a (since it is to the left of b). But then $\{y\} \leftrightarrow \{a\}$, since $\overleftarrow{\mathrm{w}}_y$ passes above a and $\overrightarrow{\mathrm{w}}_y$ is directed downwards. □

Theorem 2.4. *Let $S_1 = \{a, b, c\}$ and S_2 be two triplets of points in the plane and let $\alpha = 2\pi/3$. Assume that the wedges of S_1 and, independently, of S_2 are oriented according to the proof of Claim 2.1. Then, the induced graph $G_{S_1 \cup S_2}$ is connected.*

Proof. If one (or both) of the induced graphs G_{S_1}, G_{S_2} is a clique, then, by Lemma 2.3, we are done. Assume therefore that none of them is a clique. Let c' be the intersection point of $\overleftarrow{\mathrm{w}}_a$ and $\overleftarrow{\mathrm{w}}_c$, and consider the wedge $\mathrm{w}_{c'}$ of orientation $\theta(\mathrm{w}_{c'}) = \theta(\mathrm{w}_c)$ and apex c'. The graph induced by $\{a, b, c'\}$ is a clique, and therefore, by Lemma 2.3, $\{a, b, c'\} \leftrightarrow \{S_2\}$. If $\{a, b\} \leftrightarrow \{S_2\}$, then we are done, so assume that $\{c'\} \leftrightarrow \{S_2\}$. Let x be a point of S_2 such that $\{x\} \leftrightarrow \{c'\}$, and assume that $\{x\} \not\leftrightarrow \{c\}$ (otherwise we are done). Then, x lies above l and $\overleftarrow{\mathrm{w}}_x$ intersects $\overline{cc'}$. We distinguish between three cases, as in the proof of Lemma 2.3: (i) $\overrightarrow{\mathrm{w}}_x$ intersects $\overline{bc'}$, (ii) $\overrightarrow{\mathrm{w}}_x$ intersects l to the left of b, and (iii) $\overrightarrow{\mathrm{w}}_x$ does not intersect l. As mentioned in the proof of Lemma 2.3, our arguments there for Case (i) and Cases (ii)(1) and (ii)(2) do not use the extra assumption that G_{S_2} is a clique. Therefore, we can reuse them here. It remains to show that $\{S_1\} \leftrightarrow \{S_2\}$ in Cases (ii)(3) and (iii).

Case (ii)(3): w_y covers exactly one point of S_2, namely, c. We know that either $\theta(\mathrm{w}_y) \in R_1$ or $\theta(\mathrm{w}_y) \in R_3$. In the latter case, w_y must also cover b, which is impossible. In the former case, if y is above l, then $\{y\} \leftrightarrow \{c\}$, so y is necessarily below l. Let z be the remaining point. Then, $\theta(\mathrm{w}_z) \in R_3$. At this point, we would like to show, as in the proof of Lemma 2.3, that $\{z\} \leftrightarrow \{a, b\}$. However, we cannot assume now that $\{y\} \leftrightarrow \{z\}$. So, we first prove that $\{y\} \leftrightarrow \{z\}$, by proving that $\{x\} \not\leftrightarrow \{y\}$, and then we proceed as in the proof of Lemma 2.3.

Thus, our goal now is to prove that $\{x\} \not\leftrightarrow \{y\}$. Let p be the midpoint of $\overline{bc}$, and let a' be the projection of a onto l. According to the construction in the proof of Claim 2.1, a' lies somewhere between p and c (not including c). Let o be the intersection point of $\overleftarrow{\mathrm{w}}_x$ and l. We know that o is somewhere between c' and c (not including c). Finally, let t be the intersection point of $bis(\mathrm{w}_x)$ and l (see Figure 7(a)). We show that t lies to the left of p and therefore also to the left of a'. If t is to the left of b (or $t = b$), then this is clear. Assume therefore that t is to the right of b, and consider the two triangles $\triangle xto$ and $\triangle xbt$. Recall first that x is above $\overleftarrow{\mathrm{w}}_b$ and notice that it is below $bis(\mathrm{w}_c)$ (since, if x were above $bis(\mathrm{w}_c)$, then $\{x\} \leftrightarrow \{c\}$). Therefore $\angle xbt > 60$ and the projection of x onto l lies to the left of p. Now, in $\triangle xto$, $\angle xot \leq 60$ and $\angle txo = 60$, and therefore $|xt| \leq |to|$. And, in $\triangle xbt$, $\angle bxt < 60$ and $\angle xbt > 60$, and therefore $|bt| < |xt|$. Together, we get

that $|bt| < |to| < |tc|$, so t lies to the left of p and therefore to the left of a'.

Since the projection of x onto l lies to the left of p and so does t, we have that a lies to the left of $bis(\mathrm{w}_x)$. Now, if $\{x\} \leftrightarrow \{y\}$, then y must lie to the right of $bis(\mathrm{w}_x)$ and therefore cover a, which is impossible. We conclude that $\{x\} \not\leftrightarrow \{y\}$, and therefore $\{y\} \leftrightarrow \{z\}$ (and $\{x\} \leftrightarrow \{z\}$).

From this point, we continue as in the proof of Lemma 2.3. Notice that $\overleftarrow{\mathrm{w}}_y$ separates between a and c and between b and c, since $\theta(\mathrm{w}_y) \in R_1$ and w_y covers only c. Since $\{y\} \leftrightarrow \{z\}$, we know that z lies to the right of $bis(\mathrm{w}_y)$. Clearly, a and b lie to the left of $\overrightarrow{\mathrm{w}}_z$ (whose orientation is in R_2), and to the right of $\overleftarrow{\mathrm{w}}_z$ (whose orientation is in R_4). In other words, w_z covers both a and b. Notice also that $z \notin R_c$, since $bis(\mathrm{w}_y)$ (whose orientation is in R_1) intersects l to the right of b, and z lies to the right of $bis(\mathrm{w}_y)$. Therefore, either w_a or w_b (or both) covers z. We conclude that $\{z\} \leftrightarrow \{a, b\}$.

Case (iii): $\overrightarrow{\mathrm{w}}_x$ does not intersect l, implying that $\theta(\overrightarrow{\mathrm{w}}_x) < 180$. Notice that in this case $b \in \mathrm{w}_x$, so we assume that $x \notin \mathrm{w}_b$, implying that $\theta(\overleftarrow{\mathrm{w}}_x) > 240$. It follows that $\theta(\mathrm{w}_x) \in R_4$, $\theta(\overleftarrow{\mathrm{w}}_x) \in R_5$, $\theta(\overrightarrow{\mathrm{w}}_x) \in R_3$, and $\theta(\widetilde{\mathrm{w}}_x) \in R_1$. Notice also that $bis(\mathrm{w}_x)$, whose orientation is in R_4, intersects l to the left of b.

Let y be the point of S_2 such that $\theta(\overleftarrow{\mathrm{w}}_y) \in R_3$ and $\theta(\overrightarrow{\mathrm{w}}_y) \in R_1$, and let z be the point of S_2 such that $\theta(\overleftarrow{\mathrm{w}}_z) \in R_1$ and $\theta(\overrightarrow{\mathrm{w}}_z) \in R_5$. Notice that for $\overleftarrow{\mathrm{w}}_x$ to intersect l to the right of c', x must lie above $l(\overleftarrow{\mathrm{w}}_a)$, and, therefore, w_x covers a.

We first show that if $\{x\} \leftrightarrow \{z\}$, then $\{S_1\} \leftrightarrow \{S_2\}$. Indeed, if $\{x\} \leftrightarrow \{z\}$, then z must lie to the right of $bis(\mathrm{w}_x)$. If z is above l, then $\{z\} \leftrightarrow \{c\}$. Assume, therefore, that z is below l. Notice that $\overleftarrow{\mathrm{w}}_z$ intersects $\overleftarrow{\mathrm{w}}_b$ at a point above x, implying that $\overleftarrow{\mathrm{w}}_z$ passes above a. Moreover, $\overrightarrow{\mathrm{w}}_z$ passes below a, since it is directed downwards. It follows that w_z covers a. But, $z \in \mathrm{w}_a$, since z lies to the right of $bis(\mathrm{w}_x)$, which intersects l to the left of b. We conclude that $\{z\} \leftrightarrow \{a\}$.

Next, we address the most difficult case, in which $\{x\} \not\leftrightarrow \{z\}$. If $\{x\} \not\leftrightarrow \{z\}$, then necessarily y is connected to both x and z. Notice that z must lie below $\overleftarrow{\mathrm{w}}_x$. Also, if it is above l, then $\{z\} \leftrightarrow \{c\}$. Assume, therefore, that z is below l. Since w_y's rays are directed upwards and $\{y\} \leftrightarrow \{z\}$, we know that y is below z and therefore also below l. According to the construction in the proof of Claim 2.1, either x or z lies on $bis(\mathrm{w}_y)$, and the angle at this point in $\triangle xyz$ does not exceed the angle at the other point. It follows that the point that lies on $bis(\mathrm{w}_y)$ is necessarily x. Since, if it were z, then $\angle yzx \geq 120$, as it contains w_z.

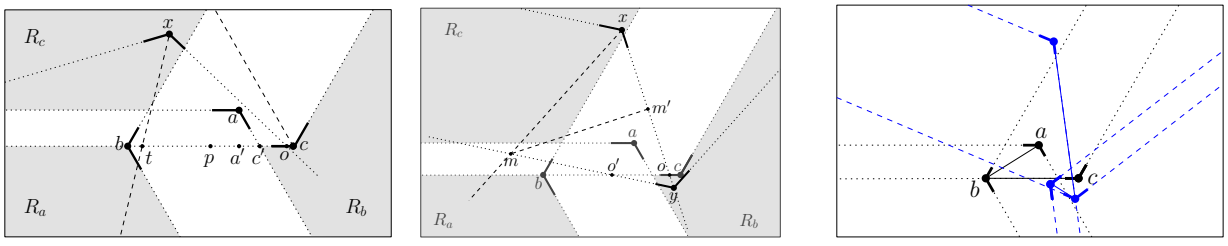

Figure 7: (a) t lies to the left of a'. (b) x lies on $bis(\mathrm{w}_y)$. (c) Each of the triplets induces a connected graph and covers the plane, but the graph of their union is not connected.

The case where x lies on $bis(\mathrm{w}_y)$ is also impossible, as we show below (see Figure 7(b)). If $y \notin \mathrm{w}_b$, then $\{y\} \leftrightarrow \{a\}$, since $\overleftarrow{\mathrm{w}}_y$ is below $l(\overrightarrow{\mathrm{w}}_b)$ and $\overrightarrow{\mathrm{w}}_y$ is above $l(\overleftarrow{\mathrm{w}}_x)$. Assume, therefore, that $y \in \mathrm{w}_b$ but $\{y\} \not\leftrightarrow \{b\}$. Let m be the intersection point of $\overleftarrow{\mathrm{w}}_y$ and $bis(\mathrm{w}_x)$. Then, m is above l

(since otherwise $\{y\} \leftrightarrow \{b\}$). Notice that $\triangle xmy$ is equilateral, and consider the bisector of $\angle xmy$. Let m' be the intersection point of this bisector and side $\overline{xy}$. Then, $\overline{mm'}$ is the perpendicular bisector of $\overline{xy}$.

Next, we show that m' lies above l. Let o be the intersection point of $\overline{xy}$ and l, and let o' the intersection point of $\overline{my}$ and l. We show that $|\overline{yo}| < |\overline{xo}|$, implying that m' is somewhere between o and x and thus above l. Consider $\triangle yoo'$. Since $\theta(\overleftarrow{\mathrm{W}_y}) \in R_3$, we know that $\angle yo'o < 60$. But $\angle oyo' = 60$, so we get that $|\overline{oy}| < |\overline{oo'}|$. Now, consider $\triangle xbo$. $\angle xbo > 60$ and $\angle bxo < 60$, and therefore $|\overline{ox}| > |\overline{ob}|$. It follows that $|\overline{oy}| < |\overline{oo'}| < |\overline{ob}| < |\overline{ox}|$.

Since all its corners lie above l, $\triangle mm'x$ is above l. Since $\{y\} \leftrightarrow \{z\}$ and z is below l, we have that $z \in \triangle yo'o \subseteq \triangle ymm'$, and therefore z is closer to y than to x – contradiction the construction of Claim 2.1. □

Remark. Theorem 2.4 above proves that when the wedges of each of the triplets are oriented, independently, according to the construction of Claim 2.1, then there is always an edge between the two triplets. This is not necessarily true for other constructions with similar properties. For example, the wedges of each of the triplets in Figure 7(c) form a connected graph and cover the plane, but there is no edge between the triplets.

3 Approximating the α-MST

Let P be a set of n points in the plane. In this section we consider the problem of computing an α-MST of P, for $\alpha = \pi, 2\pi/3, \pi/2$. For each of these angles, we devise a constant-factor approximation algorithm. The approximation ratios are actually with respect to the weight of a Euclidean MST, which is a lower bound for the weight of an α-MST, for any α.

Consider the TSP tour $\Pi = e_0, e_1, \ldots, e_{n-1}$ obtained by applying the standard 2-approximation algorithm for metric TSP. This algorithm first duplicates the edges of a MST to obtain an Eulerian tour, and then transforms the Eulerian tour into a TSP tour by introducing shortcuts. Thus, $wt(\Pi) \leq 2wt(\text{MST})$. Each of our approximation algorithms below begins by constructing Π. It then constructs, using Π, a connected α-graph, i.e., a graph in which, for each node p, the angle spanned by the edges adjacent to p is at most α. Finally, it construct an α-ST from the α-graph, whose weight is bounded by $c \cdot wt(\Pi)$, for some constant $c = c(\alpha)$, and thus is a $2c$-approximation of an α-MST.

$\alpha = \pi$. Observe that any graph of maximum degree two is a π-graph. In particular, Π is a π-graph, and, by removing an arbitrary edge, we obtain a π-ST of weight at most $2wt(\pi\text{-MST})$.

$\alpha = 2\pi/3$. Assume, for convenience, that $n = 3m$, for some integer m. We partition P into m triplets, by traversing Π from an arbitrary point $p \in P$. That is, each of the triplets consists of three consecutive points along Π. Orient the wedges of each triplet, independently, according to Claim 2.1. By Theorem 2.4, the graph induced by P, denoted here G_α (instead of G_P), is connected. In particular, for any two consecutive triplets t, t' along Π, there exists an edge of the graph between a point of t and a point of t'.

Next, we construct a $2\pi/3$-ST, T, and show that $wt(T) \leq 6 \cdot wt(2\pi/3\text{-MST})$. Initially, T has no edges. For each of the m triplets t, add to T any two edges (of the at least two edges) of G_α connecting between pairs of points of t. We call these edges *inner-edges*. Next, for each of the m

pairs of consecutive triplets t, t' along Π (except for the pair consisting of the 'last' triplet and the 'first' triplet), add to T any edge (of the at least one edge) of G_α connecting between a point of t and a point of t'. We call these edges *connecting-edges*. T is connected and has $2n/3$ inner-edges and $n/3 - 1$ connecting-edges, thus the total number of edges is $n - 1$, and T is a $2\pi/3$-ST.

We now bound the weight of T. By the triangle inequality, the weight of an edge (u, v) of T does not exceed the weight of the shorter path (in terms of number of edges) in Π between u and v. We charge the weight of this path for the edge (u, v). Each edge of Π between two points of the same triplet t is charged at most four times. Twice for the two inner-edges chosen for t, and twice for the two connecting-edges that connect t to its two adjacent triplets along Π. Each edge of Π between two consecutive triplets t, t' (except for the edge between the last and first) is charged only once for the corresponding connecting-edge of T. Thus, each edge of Π is charged at most four times, and $wt(T) = \Sigma_{e\in T}|e| \leq 4\Sigma_{e\in\Pi}|e| = 4 \cdot wt(\Pi) \leq 8 \cdot wt(\text{MST}) \leq 8 \cdot wt(2\pi/3\text{-MST})$.

Next, we improve the approximation ratio. Observe, that there are three possible ways to partition Π into m triplets. In other words, the set of edges of Π connecting between the triplets can be either E_0, E_1, or E_2, where $E_j = \{e_i \in E : i = (j \mod 3)\}$, for $0 \leq j \leq 2$. By the pigeon hole principle, the weight of one of these sets, say E_2, is at least $\frac{1}{3} \cdot wt(\Pi)$. We partition Π into triplets, such that the set of edges connecting between the triples is E_2. Now, each of the edges of E_2 (except e_{n-1}) is charged exactly once, and each of the edges of $E_0 \cup E_1$ is charged at most four times. Thus, $wt(T) \leq wt(E_2) + 4(wt(E_0) + wt(E_1)) = wt(\Pi) + 3(wt(E_0) + wt(E_1)) \leq wt(\Pi) + 3 \cdot \frac{2}{3}wt(\Pi) = 3 \cdot wt(\Pi) \leq 6 \cdot wt(\text{MST}) \leq 6 \cdot wt(2\pi/3\text{-MST})$.

$\alpha = \pi/2$. Assume, for convenience, that $n = 8m$, for some integer m. Our construction for $\alpha = \pi/2$ is similar to the one for $\alpha = 2\pi/3$, but slightly more complicated. It is based on a basic gadget described by Aschner et al. [5] for a set S of four points, indicating the locations of four $\pi/2$-wedges. This gadget is presented as the proof for the claim that one can orient the wedges of S, such that the induced graph is connected, and the wedges of S cover the plane. Unfortunately, we cannot claim that two quadruplets, whose wedges are oriented independently, are connected. However, if they are separable by a line, then they are connected, see [5].

We use this latter claim in our construction. We partition the tour Π into m sections, each consisting of 8 consecutive points along Π. Then, we partition each of the sections into two quadruplets, a left quadruplet consisting of the 4 leftmost points of the section and a right quadruplet consisting of the 4 rightmost points. (Notice that the points of a quadruplet are not necessarily consecutive along Π.) Thus, in each section, the two quadruplets are separable by a (vertical) line. Now, orient the wedges of each quadruplet, independently, such that their induced graph is connected and the wedges cover the plane. Let G_α be the graph induced by P. Observe that G_α is connected, since, for any two consecutive sections, there exists two quadruplets, one from each section, that are separable by a (vertical) line and thus connected.

Next, we construct the tree T from G_α. We distinguish between three types of edges. The first type are the *inner-edges*, which connect between points of the same quadruplet. For each quadruplet, we pick three such edges that make the quadruplet connected. The second type are the *q-connecting-edges*, which connect between quadruplets of the same section. For each section, we pick one such edge. The third type are the *s-connecting-edges*, which connect between consecutive sections along Π. For each pair of consecutive sections along Π (except for the pair consisting of the last and first sections), we pick one such edge. Notice that T is a $\pi/2$-ST, since it is connected and it has $n - 1$ edges, i.e., $3n/4$ inner-edges, $n/8$ q-connecting-edges, and $n/8 - 1$ s-connecting-edges.

We compute the approximation ratio by charging the edges of Π. Each edge of Π either connects between points of the same section, or between points of consecutive sections. An edge of the former kind is charged at most nine times. Since for a section, we have six inner-edges, one q-connecting-edge, and two s-connecting-edges. An edge of the latter kind is charged only once.

As for $\alpha = 2\pi/3$, we can choose the subset of edges of Π that connect between consecutive sections, so that its weight is at least $\frac{1}{8} \cdot wt(\Pi)$. Let E_7 denote this subset. Then, $wt(T) \leq wt(E_7) + 9 \cdot wt(E \setminus E_7) \leq wt(\Pi) + 8 \cdot wt(E \setminus E_7) \leq wt(\Pi) + 8 \cdot \frac{7}{8} wt(\Pi) = 8 \cdot wt(\Pi) \leq 16 \cdot wt(\text{MST}) \leq 16 \cdot wt(\pi/2\text{-MST})$.

The following theorem summarizes the results of this section.

Theorem 3.1. *Let P be a set of points in the plane. Then, one can construct (i) a π-ST of weight at most $2 \cdot wt(\pi\text{-MST})$, (ii) a $2\pi/3$-ST of weight at most $6 \cdot wt(2\pi/3\text{-MST})$, and (iii) a $\pi/2$-ST of weight at most $16 \cdot wt(\pi/2\text{-MST})$.*

Remark. As mentioned, the approximation ratios above are with respect to $wt(\text{MST})$, which is a lower bound for $wt(\alpha\text{-MST})$. It is possible that by comparing the weight of the constructed α-ST with that of an α-MST, one can get better ratios, but it is not clear how to do so. Moreover, it is easy to see that, for $\alpha \in [60, 180)$, 2 is a lower bound on the ratio with respect to a MST, e.g., consider n points on a line. And, for $\alpha \in [180, 240)$, $\frac{2+\sqrt{3}}{3} \approx 1.244$ is a lower bound on the ratio, e.g., consider 3 points at the corners of an equilateral triangle and a fourth point at the center of the circle passing through them. Finally, for $\alpha \in [60, 90)$, it is easy to give an example where $wt(\alpha\text{-MST})/wt(\text{MST}) \to n-1$. Therefore, any construction algorithm for an angle α in this range, should be analyzed with respect to $wt(\alpha\text{-MST})$.

4 Constant range hop-spanner for $\alpha = 2\pi/3$

In this section we apply Theorem 2.4 to obtain a solution to a problem that arises in wireless communication networks. Let P be a set of n points in the plane, where each point in P represents a transceiver equipped with an omni-directional antenna. The coverage region of p's antenna is modeled by a disk centered at p, and assume that all disks are of radius 1. Then, the resulting communication graph is the *unit disk graph* of P, denoted $\text{UDG}(P)$. (I.e, there is an edge between points p and q if the distance between them is at most 1.) As mentioned in the introduction, directional antennas have some advantages over omni-directional antennas and are gaining popularity. The coverage region of a directional antenna of angle α is modeled by a circular sector of angle α.

Assume that $\text{UDG}(P)$ is connected. Before stating our problem, we need the following definition. A graph $G = (P, E)$ is a *c-hop-spanner* of $\text{UDG}(P)$, for some constant c, if for any two points $p, q \in P$, the minimum number of hops between p and q in G is at most c times this number in $\text{UDG}(P)$. That is, for each edge $e = (p, q)$ in $\text{UDG}(P)$, there exists a path in G between p and q consisting of at most c edges. Assume now that one replaces each of the omni-directional antennas by a directional antenna of angle $2\pi/3$. We address the following *Antenna Conversion* problem: Orient the directional antennas and fix a range $\delta = O(1)$, such that the resulting (symmetric) communication graph is a c-hop-spanner of $\text{UDG}(P)$, for some constant c. I.e., construct a $2\pi/3$-graph, such that the length of its edges is bounded by δ and it is a c-hop-spanner of $\text{UDG}(P)$.

We show how to construct such a graph with $\delta = 7$ and $c = 6$, in $O(n \log n)$ time. We first partition the points of P into connected components (of $\text{UDG}(P)$) of size at most three. This is

done greedily. Set $Q = P$. As long as $Q \neq \emptyset$, perform the following step, which finds the next component C. Pick any point $a \in Q$, add it to C (which is initially empty), and remove it from Q. Now, if $Q \neq \emptyset$ and there exists a point in Q whose distance from a is at most 1, then pick any such point $b \in Q$, add it to C, and remove it from Q. Finally, if $Q \neq \emptyset$ and there exists a point in Q whose distance to either a or b (or both) is at most 1, then pick any such point $c \in Q$, add it to C, and remove it from Q.

Claim 4.1. *Let C be a connected component of size one or two. Then, each of the neighbors of C in $\text{UDG}(P)$ belongs to a component of size three.*

Proof. Assume that one of the neighbors of C belongs to a component C' of size one or two, i.e., there exists an edge of $\text{UDG}(P)$ between a point in C and a point in C'. Moreover, assume, e.g., that C was found before C'. Then, in the iteration in which C was found, we would have found a larger component, i.e., with at least one additional point. □

Now, consider the connected components that were found. We first orient the wedges of each connected component of size exactly three, independently, according to the proof of Claim 2.1. Next, for each connected component C of size one or two, let C' be any connected component of size exactly three, such that C has a neighbor in C'. Recall that the wedges of C' cover the plane. We orient each of the wedges of C (alternatively, the single wedge of C) towards the wedge of C' that covers it. Observe that if the length of the edges is not limited, then the $2\pi/3$-graph, G_α, that is induced by the wedges of P is connected. Moreover, it is easy to verify that G_α is a c-hop-spanner, for $c = 5$. However, our goal is to limit the length of the edges without increasing c by much.

Let C be a component of size one or two. Then, the edge of G_α connecting between C and C', where C' is the component of size three to which C was connected, is of length at most 4. Moreover, consider any two components of size three C' and C'', such that C has a neighbor both in C' and in C''. Then, the edge of G_α connecting between C' and C'' is of length at most 7. Finally, the edge of G_α connecting between two neighboring components of size three is of length at most 5. Therefore, one can drop all edges of length greater than 7 from G_α, without disconnecting it.

Next, we show that the resulting graph G is a 6-hop spanner. Let (p, q) be an edge of $\text{UDG}(P)$. We show that the number of hops between p and q in G is at most 6. If p, q are in the same component of size three, then, clearly, the path between them consists of at most 2 edges. If p, q are in the same component of size two, then, the path between them passes through a single component of size three, and thus consists of at most 4 edges. If p, q are in different components, then, by Claim 4.1, at least one of them, say p, is in a component of size three. If q is also in a component of size three, then the path from p to q consists of at most 5 edges. Otherwise, the path between them goes from q to some component C' of size three (which is not necessarily p's component) and from there to p's component, and thus consists of at most 6 edges.

The following theorem summarizes the result of this section.

Theorem 4.2. *Let P be a set of points in the plane and assume that $\text{UDG}(P)$ is connected. Let $\alpha = 2\pi/3$. Then, one can construct, in $O(n \log n)$ time, a 6-hop-spanner of $\text{UDG}(P)$, in which each edge is of length at most 7.*

Running time. It is possible to implement the algorithm described above in $O(n \log n)$ time, using a data structure presented by Efrat et al. [12]. This data structure is designed for a given set

Q of n points in the plane. It supports queries of the following form: Given a query point a, return a point $q \in Q$ whose distance from a is at most 1, and also delete it from Q (if requested to do so). The data structure can be constructed in $O(n \log n)$ time, and a query (including deletion if needed) can be answered in amortized $O(\log n)$ time. We use it in both phases of the algorithm. For the first phase, in which P is partitioned into connected components of $\text{UDG}(P)$ of size at most three, we construct the data structure for the set P. Now, finding a single component requires at most three queries (plus deletions), and, since there are $O(n)$ components, the total running time of this phase is $O(n \log n)$. In the second phase, we orient the wedges of each component. Orienting the wedges of the components of size three can be done in $O(1)$ time per component. For the components of size one or two, we construct the data structure for the subset of P consisting of all points belonging to components of size three. Now, by Claim 4.1, orienting the wedge of a component of size one (alternatively, the wedges of a component of size two) can be done in amortized $O(\log n)$ time. For a component of size one, we perform a single query (without deletion) in the data structure, and, for a component of size two, we perform one or two queries (without deletion), depending on whether the query with the first point is successful or not. We conclude that the overall running time of the algorithm for constructing G is $O(n \log n)$.

5 NP-hardness

We prove that the problem of computing an α-MST, for $\alpha = \pi$ and $\alpha = 2\pi/3$, is NP-hard.

5.1 $\alpha = \pi$

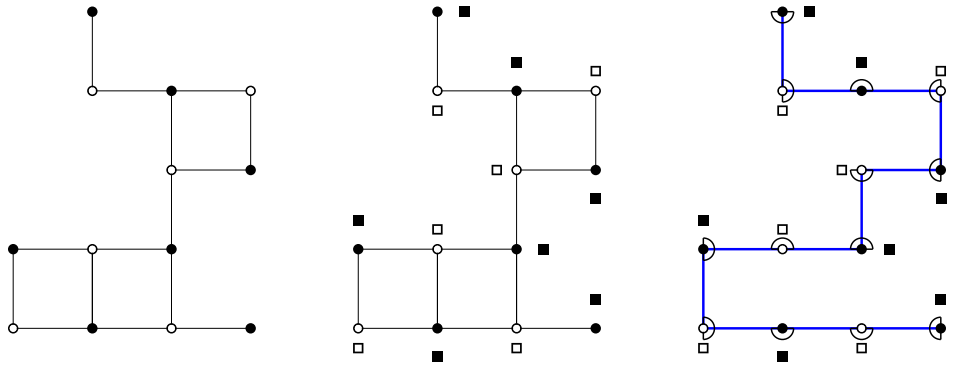

Figure 8: (a) A square grid graph $G = (V, E)$ consisting of 13 vertices and a 2-coloring of the graph. (b) For each point $v \in V$, we add a corresponding point q_v (denoted by a square). (c) A Hamiltonian path of G. The semi-circles indicate the orientation of the wedges at the points of V.

We describe a reduction from the problem of finding a Hamiltonian path in square grid graphs of degree at most 3. This problem was shown to be NP-hard by Itai et al. [14]. A *square grid graph* is a graph whose vertices correspond to points in the plane with integer coordinates, and there is an edge between two vertices if the distance between their corresponding points is 1. Let $G = (V, E)$, $|V| = n$, be a square grid graph of degree at most 3. Our reduction is very similar to

the one in [19]. Since every square grid graph is bipartite, one can color G with two colors, say, black and white (see Figure 8(a)). Let $c(v)$ denote v's color in some 2-coloring of G, and let n_1 and n_2, $n_1 + n_2 = n$, be the number of black and white points (i.e., vertices), respectively. We add a set $Q(V)$ of n points as follows. For each point $v \in V$, consider any edge e of the *complete* grid graph that is adjacent to v and is missing in G. We place a point q_v on e, such that $|vq_v| = 1/4$, if $c(v)$ is black, and $|vq_v| = 1/5$, if $c(v)$ is white (see Figure 8(b)). Notice that the distance from q_v to any point in $V \cup Q(V) \setminus \{v, q_v\}$ is greater than 1. Therefore, any MST of the point set $V \cup Q(V)$ contains the n edges $(v, q(v))$, $v \in V$, and $n-1$ edges from E. Its weight is $L = n-1+n_1/4+n_2/5$.

Lemma 5.1. *G has a Hamiltonian path if and only if $V \cup Q(V)$ has a π-MST of weight L.*

Proof. Assume G has a Hamiltonian path, then its weight is $n-1$, since it consists of $n-1$ edges and the weight of each edge in G is 1. By adding an edge between each point $v \in V$ and its corresponding point q_v, we obtain a tree, T, of weight L. Notice that the degree of each point in T is at most 3. Moreover, T is a π-ST of $V \cup Q(V)$, since, at each vertex of T, one can place a π-wedge that covers all its neighbors (see Figure 8(c)). Finally, the weight of T is L, and, therefore, T is a π-MST of $V \cup Q(V)$.

Assume now that $V \cup Q(V)$ has a π-MST, T, of weight L. Then, T must contain the n edges (v, q_v), $v \in V$, plus $n-1$ edges from E. Moreover, the maximum degree in T is at most 3 (since, a π-wedge can cover at most 3 orthogonal directions). We conclude that G has a spanning tree of maximum degree at most 2 of weight $n-1$. That is, G has a Hamiltonian path. □

5.2 $\alpha = 2\pi/3$

We describe a reduction from the problem of finding a Hamiltonian *path* in hexagonal grid graphs. Consider a tiling of the plane with regular hexagons of side length 1. The vertex set of an *hexagonal grid graph* is a subset of the vertices of the tiling, and there is an edge between two vertices of the graph if the distance between them is 1 (see Figure 9). The problem of finding a Hamiltonian *cycle* in such graphs was shown to be NP-hard by Arkin et al. [3]. We first show that the path version is also NP-hard.

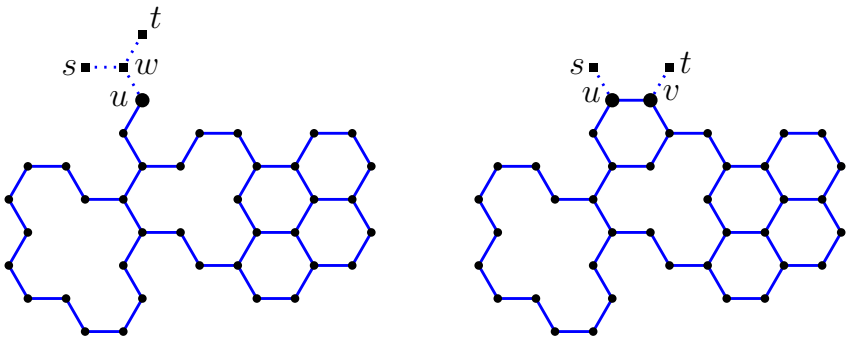

(a) (b)

Figure 9: (a) If $deg(u) = 1$, then we add three points. (b) If $deg(u) = 2$, then we add two points.

Let $G = (P, E)$ be an hexagonal grid graph, and let u be the highest point in P. (If there are several highest points, then pick the leftmost among them.). Notice that since u is the highest point, its degree cannot be 3, i.e., $deg(u) \leq 2$. Moreover, if $deg(u) = 2$, then one of the edges

adjacent to u must be horizontal. We construct another hexagonal grid graph, G', by adding at most three points to P, depending on u's degree in G. If $deg(u) = 0$, then $G' = G$. If $deg(u) = 1$, then we add the points, s, t, and w to P, as in Figure 9(a). The only edges that are formed due to this addition are (s, w), (t, w), and (u, w). Finally, if $deg(u) = 2$, then we add the points s and t to P, as in Figure 9(b). The only edges that are formed due to this addition are (s, u) and (t, v), where v is the horizontal neighbor of u.

Lemma 5.2. *G contains a Hamiltonian cycle if and only if G' contains a Hamiltonian path.*

Proof. Assume that G contains a Hamiltonian cycle. Then, $deg(u) = 2$, and (u, v) is an edge of this cycle. By dropping the edge (u, v), we obtain a Hamiltonian path between u and v in G. And, by adding the edges (s, u) and (t, v) to this path, we obtain a Hamiltonian path between s and t in G'.

Assume now that G' contains a Hamiltonian path. We claim that this is possible only if G' was obtained by adding two points to P. Indeed, if $G' = G$, then $deg(u) = 0$, and G' cannot contain a Hamiltonian path. And, if G' was obtained by adding three points to P, then, since $deg(s) = deg(t) = 1$ and s and t have a common neighbor, namely, w, G' cannot contain a Hamiltonian path. So, consider the graph G' that is obtained by adding the points s and t to P (see Figure 9(b)), and recall that we are assuming that G' contains a Hamiltonian path. Since $deg(s) = deg(t) = 1$, the endpoints of this path are necessarily s and t. By dropping the edges (s, u) and (t, v) from this path, we get a Hamiltonian path Π in G, between u and v. Notice that the edge (u, v) is not in Π, and by adding it to Π, we get a Hamiltonian cycle in G. □

We are now ready to show that the problem of computing an α-MST, for $\alpha = 2\pi/3$, is NP-hard. Let $G = (P, E)$ be an hexagonal grid graph, where $|P| = n$. Since the distance between any two points in P is at least 1, the weight of a MST of P is at least $n - 1$.

Lemma 5.3. *G has a Hamiltonian path if and only if P has a $2\pi/3$-MST of weight $n - 1$.*

Proof. Assume first that G has a Hamiltonian path. Then, its weight is $n - 1$. Moreover, the angle between any two consecutive edges along the path is $2\pi/3$. Therefore, the path is also a $2\pi/3$-ST, and, since its weight is $n - 1$, it is a $2\pi/3$-MST.

Assume now that P has a $2\pi/3$-MST, T, of weight $n - 1$. Then, all the edges of T are of length exactly 1, and therefore belong also to E. It follows that the degree of any point in T is at most 2. Therefore, T is a Hamiltonian path of G. □

References

[1] E. Ackerman, T. Gelander, and R. Pinchasi. Ice-creams and wedge graphs. *Comput. Geom.: Theory & Applications*, 46(3):213–218, 2013.

[2] O. Aichholzer, T. Hackl, M. Hoffmann, C. Huemer, A. Pór, F. Santos, B. Speckmann, and B. Vogtenhuber. Maximizing maximal angles for plane straight-line graphs. *Comput. Geom.: Theory & Applications*, 46(1):17–28, 2013.

[3] E. M. Arkin, S. P. Fekete, K. Islam, H. Meijer, J. S. B. Mitchell, Y. N. Rodríguez, V. Polishchuk, D. Rappaport, and H. Xiao. Not being (super) thin or solid is hard: A study of grid Hamiltonicity. *Comput. Geom.: Theory & Applications*, 42(6-7):582–605, 2009.

[4] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45(5):753–782, 1998.

[5] R. Aschner, M. J. Katz, and G. Morgenstern. Symmetric connectivity with directional antennas. *Comput. Geom.: Theory & Applications*, 46(9):1017–1026, 2013.

[6] I. Bárány, A. Pór, and P. Valtr. Paths with no small angles. *SIAM Journal Discrete Mathematics*, 23(4):1655–1666, 2009.

[7] P. Bose, P. Carmi, M. Damian, R. Flatland, M. J. Katz, and A. Maheshwari. Switching to directional antennas with constant increase in radius and hop distance. In *Proc. 12th Algorithms and Data Structures Sympos.*, pages 134–146, 2011.

[8] I. Caragiannis, C. Kaklamanis, E. Kranakis, D. Krizanc, and A. Wiese. Communication in wireless networks with directional antennas. In *20th ACM Sympos. on Parallelism in Algorithms and Architectures*, pages 344–351, 2008.

[9] P. Carmi, M. J. Katz, Z. Lotker, and A. Rosén. Connectivity guarantees for wireless networks with directional antennas. *Comput. Geom.: Theory & Applications*, 44(9):477–485, 2011.

[10] T. M. Chan. Euclidean bounded-degree spanning tree ratios. *Discrete & Computational Geometry*, 32(2):177–194, 2004.

[11] A. Dumitrescu, J. Pach, and G. Tóth. Drawing Hamiltonian cycles with no large angles. *Electronic Journal of Combinatorics*, 19(2):P31, 2012.

[12] A. Efrat, A. Itai, and M. J. Katz. Geometry helps in bottleneck matching and related problems. *Algorithmica*, 31(1):1–28, 2001.

[13] S. P. Fekete and G. J. Woeginger. Angle-restricted tours in the plane. *Comput. Geom.: Theory & Applications*, 8:195–218, 1997.

[14] A. Itai, C. H. Papadimitriou, and J. L. Szwarcfiter. Hamilton paths in grid graphs. *SIAM Journal on Computing*, 11(4):676–686, 1982.

[15] R. Jothi and B. Raghavachari. Degree-bounded minimum spanning trees. *Discrete Applied Mathematics*, 157(5):960–970, 2009.

[16] S. Khuller, B. Raghavachari, and N. E. Young. Low-degree spanning trees of small weight. *SIAM Journal on Computing*, 25(2):355–368, 1996.

[17] E. Kranakis, D. Krizanc, and O. Morales. Maintaining connectivity in sensor networks using directional antennae. In Theoretical Aspects of Distributed Computing in Sensor Networks, Chapter 3, S. Nikoletseas and J. D. P. Rolim (Eds.), Springer.

[18] J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: a simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999.

[19] C. H. Papadimitriou and U. V. Vazirani. On two geometric problems related to the travelling salesman problem. *Journal of Algorithms*, 5(2):231–246, 1984.